Tapping the sensorimotor trajectory

Oswald Berthold
Adaptive Systems Group
Dept. of Computer Science
Humboldt-Universität zu Berlin
Unter den Linden 6, 10999 Berlin
Email: bertolos@hu-berlin.de

Verena V. Hafner
Adaptive Systems Group
Dept. of Computer Science
Humboldt-Universität zu Berlin
Unter den Linden 6, 10999 Berlin
Email: verena.hafner@hu-berlin.de

arXiv:1704.07622v1 [cs.NE] 25 Apr 2017

***Abstract*—We propose sensorimotor tappings, a new graphical technique that explicitly represents relations between the time steps of an agent's sensorimotor loop and a single training step of an adaptive model that the agent is using internally. In the simplest case this is a *relation linking two time steps*. In realistic cases these relations can extend over several time steps and over different sensory channels. The aim is to capture the footprint of information intake relative to the agent's current time step. We argue that this view allows us to make prior considerations explicit and then use them in implementations without modification once they are established.**

In the paper we introduce the problem domain, explain the basic idea, provide example tappings for standard configurations used in developmental models, and show how tappings can be applied to problems in related fields.

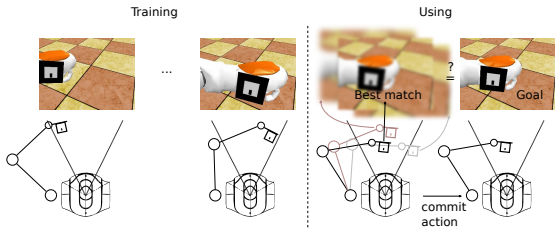

Fig. 1: Example: On the left a Nao robot trains a model to predict visual consequences from joint angle configurations through sensorimotor exploration, right: the robot uses the model to find the best matching prediction and the associated action in the predictor's input.

I. Introduction

Looking at developmental processes in higher animals reveals that they involve large amounts of learning from experience. Developmental Robotics, in agreement with Reinforcement Learning and Computational Neuroscience, is concerned with evaluating hypothetical principles of development and learning. Although the fields differ in motivation, evaluation happens to a large extent through computational models in which the principles are represented as patterns of interaction between a number of adaptive components with specific functions.

A generally accepted view is that the agent's brain contains modules which function as models of the agent's interaction with the world. The idea is that these models are used by the brain to evaluate the possible actions in "imagined space" and only commit to the most promising one in physical space. The task of a theory now is to describe precisely how a sensorimotor model is learnt from experience and how it interacts with other models present within the context of a developmental model.

We need to emphasize this distinction between different types of models. Machine learning (ML) solves the problem of how to build models from data in a problem independent form. The ML approach relies on a *preprocessing step* to transform the raw data into the required form. Using ML methods we can learn sensorimotor models of transitions in sensorimotor space up to a desired accuracy. This level of modelling provides the grounding in sensorimotor space. Developmental models need to capture the preprocessing and the interaction of the sensorimotor models that they contain. Sensorimotor models can have different roles which are determined by the relations contained in the training data. An important question is *how to map the raw sensorimotor data to sensorimotor training data* for realizing specific functions needed inside a developmental model.

Consider a simple example from robotics involving an experiment like the one shown in Figure 1. A humanoid robot is bootstrapping the ability to move one of its hands to a specified point in the visual space of the robot's camera. For simplicity the camera does not move. The robot's arm has four joints and the hand's position is determined by the configuration of all four individual joints taken together. The agent records a series of random joint angle commands and corresponding camera images. Using these data, it trains a sensorimotor model to predict the resulting camera image from the joint angles that generated it. Finally the agent exploits the trained model by feeding random commands into the model, recording the predictions and choosing the commands which result in the prediction closest to the goal. For all of this to work, the sensorimotor training data at any time step need to be composed of the image at that same time step and the command from the previous time step. In the rest of the paper we develop a graphical technique for representing and analzying such mappings from sensorimotor time steps to training steps.

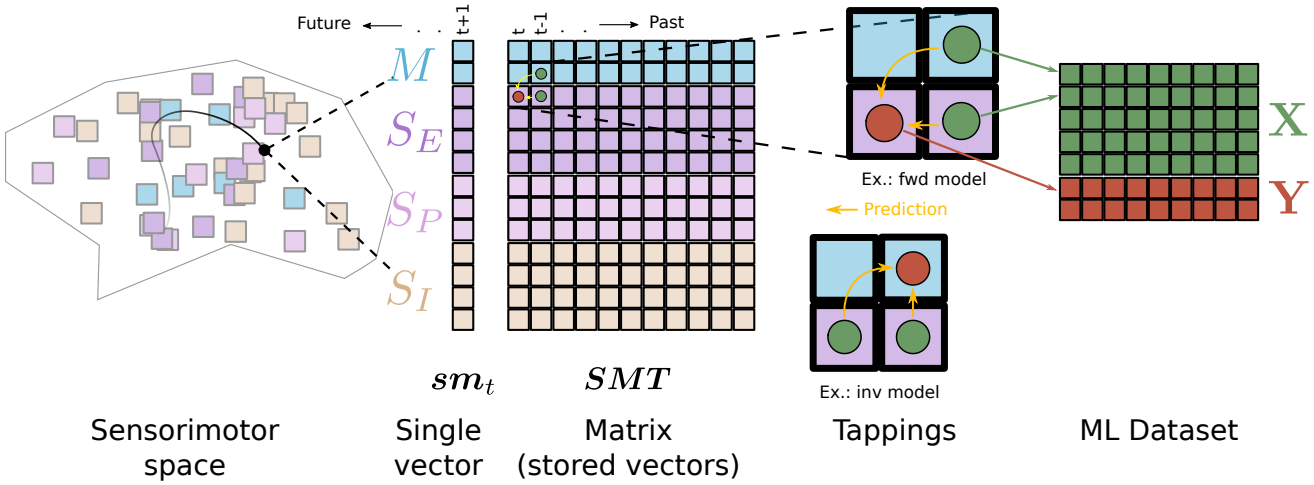

Fig. 2: The basic idea of tapping the sensorimotor trajectory. Concatenating the row vectors horizontally creates a matrix. The matrix inherits the row structure from the vector and represents time along the other axis.

A. *Related work*

The techniques most commonly used for discussing developmental models are either plain textual descriptions, equations in formal approaches, or variations of block diagrams and signal-flow diagrams on the graphical side. Nothing in general can be said about plain text descriptions. Equations and flow diagrams are usually highlighting different but fixed aspects of the model's function or the design problem. Equations are precise in representing functional dependencies and time. Block- and flow diagrams tend to emphasize the structure of interaction among the variables of interest. For analysis, these techniques are usually complemented by problem specific graphical representations of behaviour.

The visualisation and conceptualisation techniques are coming largely from electrical diagrams, signal processing and control theory. The ideas of black boxes and modular design, as they are used in block diagrams, is rooted in these fields as well. The black box view allows for general input/ouput analyses. The corresponding equations, for example discrete state space based iterative update rules, equally reflect their origin in the formalisms of these areas. A notable exception are the backup diagrams introduced by Sutton & Barto in [1]. Backup diagrams track how the instantaneous information is related to previous states and indicates how it is propagated back in time to update the relevant state in the agent's controller. These diagrams do not however differentiate sensory modalities.

The most important concept in digital signal processing for time domain manipulation is the finite impulse response filter. Such a filter's output is computed as a weighted sum over a finite number of past inputs including the current one. It is realized by *tapping* into a sliding window over inputs with fixed real-valued weights. Each weight corresponds to a fixed delay and the largest delay with a non-zero weight is the maximum time across which there can be causal interaction between input and output.

Information theory can be applied to analyze the information flow across the inputs and outputs of a black box to characterize general properties of information *shared* among these variables. Information theoretic analysis of sensorimotor relations can be made quantitative [2] and has been proposed as a way to capture high level properties of experience [3]. A number of recent proposals for modelling behaviour as a self-organizing process have suggested to measure information sharing over time [4] and to use that measure for driving behaviour [5], [6], [7]. This can be taken as evidence for the importance of predictability for guiding agent-environment interaction. The dependence of predictability on the amount of past experience going into the prediction allows to precisely determine the memory needs of a given agent.

Internal modelling approaches in developmental robotics that use prediction learning are missing a way to describe in a general and systematic manner the interaction of the embedded sensorimotor models with the information provided by the enclosing developmental model. This also holds for temporal difference learning in RL and correlational learning processes in neuroscience. Thus we see a definite need for an additional tool from which these fields, and maybe robotics and AI at large, might benefit. Our contribution besides the identification of this gap is a proposal for filling it.

II. Tappings

The sequence of steps necessary for going from sensorimotor space to the sensorimotor models's input / output space are shown in the illustration in Figure 2 with a blow-up of the tapping transformation. In order to learn, an agent must produce information through actions that produce changes in the outside world. The information becomes available to the agent *exclusively* through its sensors [1]. A single sensory

[1]To simplify the discussion, we use *sensors* as a placeholder for any kind of internal state variable. These are either primary signals, like the actual sensors, motors or basic rewards, but also signals derived from any intermediate stage in the agent brain's flow of computation.

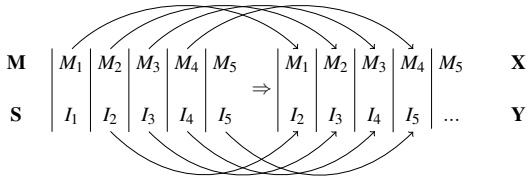

Fig. 3: Example cont.: This illustrates the remapping of time that the Nao agent needs to realize when constructing the training data, input $\boldsymbol{X}$ and labels $\boldsymbol{Y}$, from the raw sensorimotor data in this particular example.

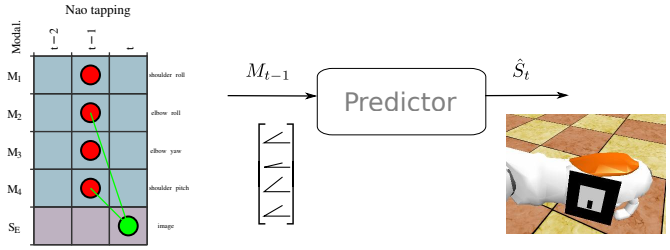

Fig. 4: Tapping for the nao example with a corresponding block diagram.

measurement is usually represented by a vector. The vector is composed of subparts that reflect the natural structure given by the types of the sensors, which correspond to *modalities* (e.g. vision or joint angles). The set of all possible vectors defines the agent's sensorimotor space. The agent's internal time creates the temporal ordering of incoming measurements [8], and storing them in this order forms a matrix. The matrix contains a numerical representation of the sequence of changes in the real world *as they are reflected* in sensorimotor space. If the agent lives in a world that is not completely observable (pretty much every world), it must squeeze out additional information by exploiting relations across time and across modalities. To do this it must tap the sensorimotor matrix using a fixed pattern attached to the current time step with the growing matrix sliding along underneath.

Before examing this in more detail let's see how this translates back into our example. A humanoid robot is bootstrapping the ability to move its hand to a given point in visual space. The agent creates information by generating random joint angle commands and recording the resulting image. For simplicity we say that the motion of the arm is finished by the time the next image is taken. Each measurement vector contains the *current* image, the previous command's consequence, and a *new* motor command about to be committed [2]. The agent records measurements over 5 time steps and trains an internal forward model to predict the image in the current time step from the previous time step's joint angle command. To do this, the agent constructs a training set for an adaptive model. The textual description is complicated although the relation is very simple as shown in Figure 3. The agent aligns the motor commands from the first to the second to last values as the input, and the images from the second to the last image as the target to construct a matrix of 5-1 training data points. The relation in this example is trivial but already contains the essential ingredients connecting two different modalties across two time steps. The prediction happens within one time step but its input can extend over any number of time steps.

Now we are prepared for a detailed account of the tapping procedure. Models of interacting agents usually have a fixed number n_s of sensors and a fixed number n_m of motors, both represented as column vectors. Each vector has additional internal structure reflecting different types and groups of sensors and actuators. In Figure 2 for example we are using four different modalities: motors, proprioception, exteroception and interoception, denoted by $\boldsymbol{M}$, $\boldsymbol{S}_\mathrm{P}$, $\boldsymbol{S}_\mathrm{E}$ and $\boldsymbol{S}_\mathrm{I}$ respectively. To obtain a general representation we concatenate all these vectors into one large column vector $\boldsymbol{sm} \in \mathbb{R}^{n_{sm}}$ with dimension n_{sm} holding all sensorimotor variables. All possible realizations of $\boldsymbol{sm}$ span the sensorimotor space $\boldsymbol{SM} = \mathbb{R}^{n_{sm}}$. In most cases the agent cannot actually realize all physical correspondences with points in $\boldsymbol{SM}$. The set of actually reachable states is called the sensorimotor manifold. The sensorimotor loop adds a discrete time t for every iteration it is going through. For every t there is a realization $\boldsymbol{sm}_t$ of the sensorimotor vector. We collect all of these and store them side by side in the sensorimotor matrix $\boldsymbol{SMT}$.

A common arrangement in a developmental model is to use a single learning algorithm inside a box for representing an internal model. Such a model is most effectively trained with supervised learning. A supervised training set consists of input data $\boldsymbol{X}$ and target data $\boldsymbol{Y}$ that constrain the functional relation $f(\boldsymbol{X}) = \boldsymbol{Y}$. The approximation task is to find parameters θ for the model $\hat{f}(\cdot, \theta)$ such that $|f - \hat{f}|$ is minimized under some measure. Prediction learning allows the agent to construct infinite supervised training data on the fly. What is missing is a systematic way to describe that construction process.

To fill this gap we propose **tappings** and define them as directed graphs on top of $\boldsymbol{SMT}$'s row and column coordinates. Each node of the graph is placed on a cell of the matrix and assigned an input or target color. The tapping is defined relative to the current time t with the data sliding along under the tapping with each time step. If $\boldsymbol{XY}$ is the full supervised training set, the tapping defines a map taking a set of $\boldsymbol{SMT}$ coordinates to a set of $\boldsymbol{XY}$ coordinates. The node coordinates on the $\boldsymbol{SMT}$ grid encode the relation prescribed by the sensorimotor model's function inside the developmental model.

Returning to our example one more time we can now write it down as the tapping shown in Figure 4. The Nao agent's sensorimotor space consists of four motor values and an image. For simplicity we leave away the details of image processing. During exploration the agent creates a sensorimotor trajectory from which it learns to predict visual consequences from joint angle commands. By design of the experiment there is a delay of one time step between committing a motor command and

[2]This ordering is only a convention but in our view the most intuitive one.

measuring the results associated with that command. This means that every time there is a new image measurement the agent creates a new training data point by concatenating the current image with the previous command and performs a single update step on the internal model.

A. Tapping degrees of freedom

The design space of tappings established it needs to be examined for the possibilities of systematic variation. The variables used for a given agent are commonly grouped during the definition of sm according to a few main properties of the what the variable represents. They are distinguished according to wether the variable encodes a sensor or motor signal, and then again according to lower level properties such as exteroceptive- (vision, hearing), proprioceptive- (joint angles, forces), or *interoceptive* signals. Interoceptive variables represent any intermediate stage of the computations in the agent's sensorimotor integration process. The term is based on the idea that the brain can deploy sensor probes to any location inside the circuits that define its own operation. Blocks of different color in a tapping indicate the group. A group whose elements all contribute to the same argument of the desired function, for example all pixels in an image, it can be reduced to a single element in the graphical representation. Groups are organized along the row axis of the matrix representing modalities.

Tappings are specified relative to the current time $t = 0$, becoming positive in the future and negative into the past. This proposal only considers discrete time and equidistant sampling with a constant Δt. The sampling interval needs to be tuned to match the processes the agent needs to observe and to provide sufficient time to compute one iteration of the sensorimotor loop. Loop iterations are unrolled along the column axis of the matrix representing agent time.

These are the principal axes of association along which a model can learn a functional relation. A single time step prediction problem requires a tapping which assigns a variable at $t-1$ to the input and the same variable at time t to the target of a model update. Doing the same along modalities captures intermodal prediction, that is, predicting sensory consequences in one modality from the state of another modality. Adding joint angle sensors the Nao agent it could learn to predict the hand position (vision) from joint angles (proprioception) in the same time step. These two proto-tappings from which all others can be composed are shown in Figure 5a and Figure 5b.

Sensorimotor models encode regularity in sensorimotor state transitions along these axes. Learning transitions along the normal forward flow of time results in a forward model. Forward models are central to the simulation theory of cognition. The theory states that an agent learning to approximate the forward transition rules to a sufficient degree of fidelity can use them to interally "simulate experience" [9].

Rearranging the direction of prediction to go backwards in time using the same set of variables used in the forward model. This allows the model to predict (infer) causes from observed effects. Technical details aside, this allows an agent to control

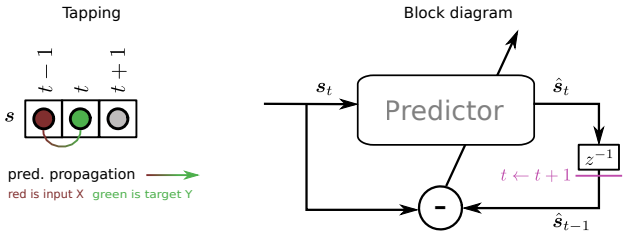

(a) Simple temporal predictor.

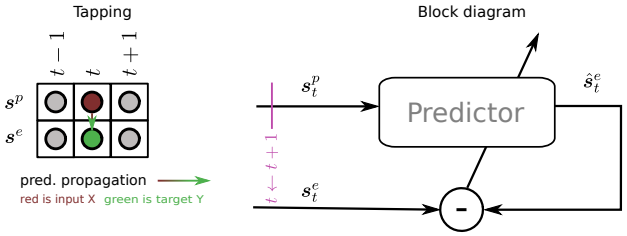

(b) Simple intermodal predictor.

Fig. 5: The two principal axes of association shown as tappings alongside with corresponding block diagrams. a) A simple temporal predictor, predicting the state one timestep ahead, and b) a simple intermodal predictor taking proprioceptive input to a exteroceptive prediction.

and change its own state by directly predicting the causes of its desired state. This translates to predicting the actions that lead to a goal [10]. This is a functional capability which can be implemented by sampling forward simulations or as a single prediction step. Direct prediction has a technical requirement on the learning algorithm. The inverse of a function is not garantueed to be a function itself but instead relate a single consequence to multiple causes. The mathematical object is called a correspondence and the problem can be solved with learning algorithms using probabilistic representations.

B. Summary

To summarize this section we highlight the main features of tappings. They provide an information centric view on developmental models. This view is independent of particular learning algorithms, and it provides an upper bound [3] on the amount of explanation a model needs to accomplish. That bound is a reference for comparing different models in terms of the fraction of maximum explanation. Tappings facilitate the design of developmental models, algorithms and their implementations by highlighting regularities in the design space and being precise and explicit about time. Analysing two important model types and their tappings shows to what extent different functional roles are determined by the input / output relations, and the learning algorithm respectively. These features all contribute to facilitate systematic exploration of developmental models.

[3] the joint entropy of all sensorimotor variables

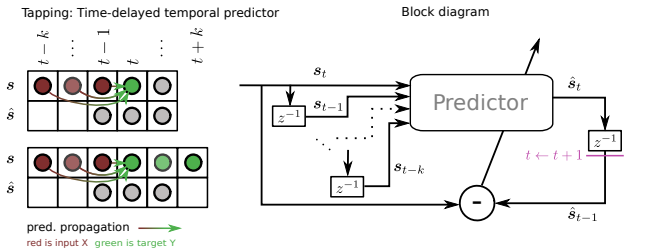

Fig. 6: The multi step predictor using a window on k past values as instantaneuous input and, in the fully symmetric case a window on $k-1$ addtional future values as the target. The time indexing has been omitted for that case for less clutter.

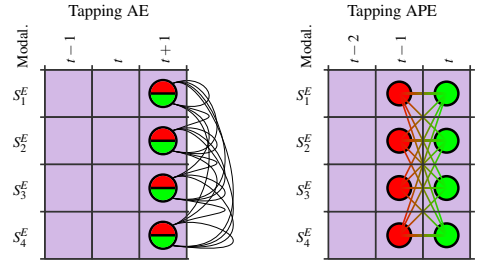

Fig. 7: Autoencoder (left) and autopredictive encoder (right). The AE's tapping is special because input and target coincide. Pulling the input and source apart over one timestep difference produces the autopredictive encoder. The prediction prior imposes additional structure on the hidden representation.

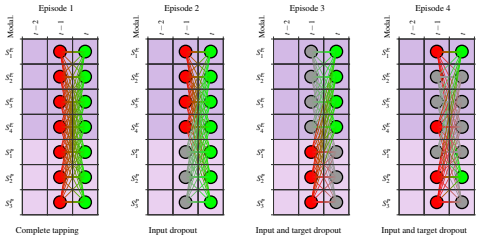

Fig. 8: This illustrates the input / output dropout technique applied to an autopredictive encoder that integrates inputs from two modalities. The technique can be implemented either by augmenting the training data or by using time variant tappings with a blocking control. Going through the diagrams from left to right they each represent a dropout configuration for one episode for starting with a complete tapping on the left, followed by input dropout, input and target dropout, and a random configuration.

III. Basic Tappings

In this section we explore tappings further by looking at some variations of the simple ones that came out of the previous section: multi step prediction, autoencoding, and autopredictive encoding. If the internal model is a feedforward map without internal memory the simple one time step predictor in Figure 5a cannot make use of additional information about the future that was presented more than one time step ago. The missing memory of the model can be replaced by using a *moving window* of size k that augments the momentary model input by including all k previous values of the variables [4]. Since tappings are moving windows, the multi time step tapping shown in Figure 6 is almost trivial, the window size being equal to the number of input taps spread uniformly into the past. The interesting bit is that iterative predictions in extended forward simulations demands better model accuracy. A reasonable shortcut towards more accuracy is to improve the prediction with imposing a long-term consistency constraint by extending the target tapping into the future (using buffering in closed-loop learning).

A special case of a predictor is the autoencoder. Its tapping is shown on the left in Figure 7. Its target output is the same as its input. In terms of the $\boldsymbol{XY}$ formulation with $\boldsymbol{X} = \boldsymbol{Y}$, the autoencoder could only consist of wires. The added value of an autoencoder comes exclusively from constraints on the intermediate representation. Like prediction learning, autoencoding is an unsupervised learning technique built with supervised learning. For many unsupervised learning tasks. If we look at the tapping we see that the information of each single variable on the input is distributed to all other variables on the output. By a simple change of the tapping we easily obtain an autopredictive encoder (APE) as the result of pulling the autoencoder's input and output taps one time step apart. The autopredictive encoder is not an established term but multiple proposals for such architectures have in fact been made [11], [12], [13]. Applying the prediction constraint on the model has been shown to increase the task-independence of latent space representations by (Lotter et al. 2015, [14]). In the tapping we see immediatly that the prediction constraint encourages the model to represent the rules of change in the hidden space. The APE tapping is shown in Figure 7. With a simple variation of the APE tapping we can implement another powerful constraint for latent space representation without changing the model. One possibility is to augment the training set by appending multiple copies of itself to the original. Before appending it, a fixed proportion of randomly chosen input entries in the copy is set to an inactive state. If taps are allowed to be blocking the tapped data, the same can be achieved by making a set of random tapping nodes to be blocking for every episode of training data that is presented. This is sketched in Figure 8. This is a version of the dropout regularization technique applied to the dataset. Like the prediction constraint this results in models that generalize better for incomplete inputs which is known as conditional inference in the Bayesian view [15], [16], [17].

[4]The moving window technique is alternatively known as moving average model, time delay neural network, delay-embedding or method of delays

IV. Application Areas

Internal modelling [18] is an important concept used in developmental robotics [19], [20], [21]. An underlying driving hypothesis is that predictive models enable *anticipatory* behaviour [22] which is more powerful than purely reactive be-

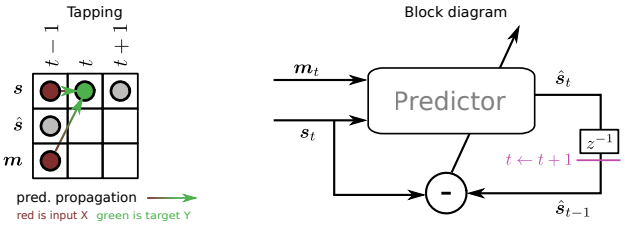

(a) Forward

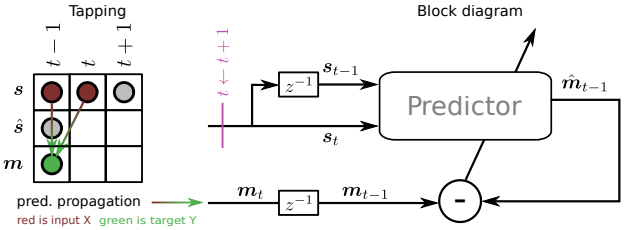

(b) Inverse

Fig. 9: Tapping a single time step forward- and inverse model pair. The model's functions are determined by different relations over the same set of variables.

haviour. From the developmental perspective this implies that some functions of a developmental model must be provided by adaptive models of the sensorimotor dynamics. Two basic functional types of internal models, forward and inverse ones, have already been introduced as examples in Figure 2 and are shown again as a pair of tappings in Figure 9. This highlights the rearrangement of the direction of prediction without a change of variables. Exploitation of adaptive models has also been described above indicating different ways of predicting and evaluating future options with forward models, or directly infering actions with inverse models.

A popular method in reinforcement learning is temporal difference learning. Temporal difference learning is a family of algorithms to approximate a prediction target with a recurrent estimate. The usual target is a *value function* which maps actions to a value. The estimate is bootstrapped by minimizing the moment-to-moment value prediction error, which is ultimately grounded in a primary reward signal. There exists extensive theory in RL that deals with the problem of integrating task-relevant information that is spread out in time, with two fundamental concepts being involved. The first one is that of *multistep methods* which take care of consequences escaping into the future. The second one are *eligibility traces* which capture causes vanishing into the past. Taken together they solve the general delayed reward problem. Depending on the parameters a corresponding tapping will be similar to the multi step predictor.

The importance of features and modalities and the information contained in their mutual relations is less developed. The concepts used in reinforcement learning can easily be remapped to internal modelling terms and vice versa, making tappings immediatly applicable to temporal difference learning problems. Looking at three basic temporal difference learning

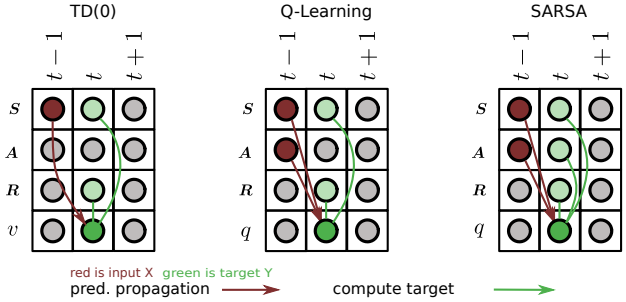

Fig. 10: Tapping temporal difference learning algorithms.

algorithms, TD(0), Q-Learning and SARSA, we see that they all approximate a target by updating from a one time step difference. TD(0)'s target is a state value function v while for Q-learning and SARSA it is a state-action value function q [1]. The update rules all follow the same general form of

$$\Delta v = \alpha(R_t + \gamma v(S')_t)$$

and the corresponding tappings are shown in Figure 10. Comparing these with the internal model tappings we see that temporal difference learning corresponds with prediction learning and that the value function is a forward model allowing us to reframe a RL problems as developmental prediction learning ones and the other way round.

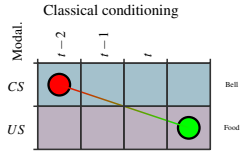

Fig. 11: Classical conditioning model requiring prediction of the unconditioned stimulus from another stimulus preceding the unconditioned one by a few time steps.

Models of classical conditioning in neuroscience, like the Rescorla-Wagner rule [23] provide links between the computational and neurobiological accounts of reinforcement learning. Such models describe how an association is learned across two modalities, the unconditioned (US) and the conditioned stimulus (CS), and across two time steps, the CS preceding the US. At the level of functional organisation the reward prediction error hypothesis of dopamine [24], [25], [26], and more recently, the predictive-brain hypothesis [27], [28] carry predictive coding ideas to the extreme. The claims are that the information flow in the brain consists of top-down predictions and bottom-up prediction errors, and that the brain is organised in adaptive layers that learn to predict signals ascending from lower layers with the help of predictions descending from higher layers. Only the prediction error is being passed on to higher layers which thereby becomes the main drive of activity. Tappings can again be applied at all of these different levels of investigation.

V. Discussion

During our presentation of tappings a few issues have come up but whose discussion was delayed for reasons of simplicity.

Models with memory like recurrent neural networks need special consideration with respect to tappings. Recurrent systems naturally retain an internal memory of past input values. Such models do not need explicit memory in their inputs and in theory only need to tap across one time step. They are building up an implicit tapping as part of their learning. Measuring the information flow across the model inputs and outputs after training with quantitative [2] or relational techniques [29] should result in what could be called effective tapping and this should be applicable to training memoryless models.

The memory issue is an example of a more general aspect about tappings. In the current state of our proposal, tappings disregard details about the learning algorithm used at the level of sensorimotor models. The same is evident in the case of inverse problems where we needed to consider learning of correspondences instead of functions. It remains to be shown how these properties could be represented in a tapping.

A final interesting issue is the problem of changes in the tapping or the sensorimotor space itself that might occur during a learning episode. When sketching the dropout constraint for the autopredictive encoder a new tapping was drawn randomly for each episode, which is a simple case of variation in time. We think it is striaghtforward to obtain a more general time dependent formulation of tappings. The tapping parameters could even be added to sensorimotor space itself as interoceptive variables. Any advanced developmental model will include an aspect of growth, just because in open ended scenarios, it is not clear a priori how much explanatory capacity is needed for a given level of accuracy in behaviour. If a learning algorithm changes its own model size during learning, new sensorimotor variables are either created or destroyed. If remains to be seen if tappings can be extended to such cases in a meaningful way.

VI. Conclusion

We have presented tappings, a novel concept for designing and analysing models of developmental learning in the field of developmental robotics and the related fields of reinforcement learning and computational neuroscience. Tappings came out of a need for capturing the detailed embedding of learning machines in the temporal and modal context of raw sensorimotor trajectories. Tappings create a particular view on the interaction between the embodiment and the functional requirements of behaviour that can help to better understand developmental learning processes, and make sensorimotor learning more efficient. They can systematically describe the relationship between supervised learning and developmental models. By ignoring computational details the tapping view highlights the information flow across models and using that we can compare a large range of models that cannot easily be compared otherwise. We showed the structural similarity of prediction learning in the developmental context and temporal difference learning in RL. Using tapping manipulations we arrived at a learning scheme based on a predictive version of the autoencoder which we think is a candidate for a general building block in developmental models.

Acknowledgment

We would like to thank Guido Schillaci for providing experimental data for the Nao example and the Adaptive Systems Group members for discussions.

References

[1] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, 1st ed. Cambridge, MA, USA: MIT Press, 1998.

[2] M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns, "Methods for quantifying the informational structure of sensory and motor data," *Neuroinformatics*, vol. 3, no. 3, pp. 243–262, 2005.

[3] G. Tononi, O. Sporns, and G. M. Edelman, "A measure for brain complexity: relating functional segregation and integration in the nervous system," *Proceedings of the National Academy of Sciences*, vol. 91, no. 11, pp. 5033–5037, 1994. [Online]. Available: http://www.pnas.org/content/91/11/5033.abstract

[4] W. Bialek and N. Tishby, "Predictive Information," *eprint arXiv:cond-mat/9902341*, Feb. 1999.

[5] R. Der and G. Martius, *The Playful Machine: Theoretical Foundation and Practical Realization of Self-Organizing Robots*, ser. Cognitive Systems Monographs. Springer Berlin Heidelberg, 2012. [Online]. Available: http://www.springer.com/us/book/9783642202520

[6] G. Martius, R. Der, and N. Ay, "Information driven self-organization of complex robotic behaviors," *PloS one*, vol. 8, no. 5, p. e63400, 2013.

[7] A. S. Klyubin, D. Polani, and C. L. Nehaniv, "Keep your options open: An information-based driving principle for sensorimotor systems," *PLOS ONE*, vol. 3, no. 12, pp. 1–14, 12 2008.

[8] A. V. Terekhov and J. K. O'Regan, "Space as an invention of active agents," *Frontiers in Robotics and AI*, vol. 3, p. 4, 2016.

[9] G. Hesslow, "The current status of the simulation theory of cognition," *Brain Research*, vol. 1428, pp. 71 – 79, 2012, the Cognitive Neuroscience of Thought. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0006899311011309

[10] M. Rolf and M. Asada, "What are goals? and if so, how many?" in *2015 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, Aug 2015, pp. 332–339.

[11] V. Michalski, R. Memisevic, and K. Konda, "Modeling deep temporal dependencies with recurrent grammar cells"","" in *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc., 2014, pp. 1925–1933. [Online]. Available: http://papers.nips.cc/paper/5549-modeling-deep-temporal-dependencies-with-recurrent-grammar-cells.pdf

[12] V. Patraucean, A. Handa, and R. Cipolla, "Spatio-temporal video autoencoder with differentiable memory," *CoRR*, vol. abs/1511.06309, 2015. [Online]. Available: http://arxiv.org/abs/1511.06309

[13] J. L. Copete, Y. Nagai, and M. Asada, "Motor development facilitates the prediction of others' actions through sensorimotor predictive learning," in *Proceedings of the 6th IEEE International Conference on Development and Learning and on Epigenetic Robotics*, 2016.

[14] W. Lotter, G. Kreiman, and D. Cox, "Unsupervised learning of visual structure using predictive generative networks," *CoRR*, vol. abs/1511.06380, 2015. [Online]. Available: http://arxiv.org/abs/1511.06380

[15] A. F. Morse, J. de Greeff, T. Belpeame, and A. Cangelosi, "Epigenetic robotics architecture (era)," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 4, pp. 325–339, Dec 2010.

[16] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, "Multimodal deep learning," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 689–696.

[17] K. Sohn, W. Shang, and H. Lee, "Improved multimodal deep learning with variation of information," in *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc., 2014, pp. 2141–2149. [Online]. Available: http://papers.nips.cc/paper/5279-improved-multimodal-deep-learning-with-variation-of-information.pdf

[18] K. J. W. Craik, *The Nature of Explanation*. Cambridge University Press, 1943.

[19] D. M. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control," *Neural Networks*, vol. 11, no. 7, pp. 1317–1329, 1998.

[20] Y. Demiris and B. Khadhouri, "Hierarchical attentive multiple models for execution and recognition of actions," *Robotics and Autonomous Systems*, vol. 54, no. 5, pp. 361 – 369, 2006. [Online]. Available: //www.sciencedirect.com/science/article/pii/S0921889006000169

[21] G. Schillaci, V. V. Hafner, and B. Lara, "Exploration behaviors, body representations, and simulation processes for the development of cognition in artificial agents," *Frontiers in Robotics and AI*, vol. 3, p. 39, 2016.

[22] R. Rosen, *Anticipatory Systems - Philosophical, Mathematical, and Methodological Foundations*, 2nd ed., ser. IFSR International Series on Systems Science and Engineering. Springer-Verlag New York, 2012, vol. 1.

[23] R. Rescorla and A. R. Wagner, "A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement," in *Classical conditioning II: Current research and theory*. Appleton-Century-Crofts, New York, 1972, p. 64–99.

[24] W. Schultz, P. Dayan, and P. R. Montague, "A neural substrate of prediction and reward," *Science*, vol. 275, no. 5306, pp. 1593–1599, 1997.

[25] P. Dayan, "Matters temporal," *Trends in Cognitive Sciences*, vol. 6, no. 3, pp. 105–106, 2002.

[26] Y. Niv, "Reinforcement learning in the brain," *Journal of Mathematical Psychology*, vol. 53, no. 3, pp. 139 – 154, 2009, special Issue: Dynamic Decision Making.

[27] K. Friston, "A free energy principle for biological systems," *Entropy*, vol. 14, no. 11, p. 2100, 2012.

[28] A. Clark, "Embodied prediction," in *Open MIND*, T. K. Metzinger and J. M. Windt, Eds. Frankfurt am Main: MIND Group, 2015, ch. 7(T).

[29] P. L. Williams and R. D. Beer, "Nonnegative Decomposition of Multivariate Information," *ArXiv e-prints*, Apr. 2010.